

The Weightless Economy in Economic Development

by

Danny Quah *
LSE Economics Department
February 1999

CENTRE FOR ECONOMIC PERFORMANCE
DISCUSSION PAPER NO. 417

This paper was produced as part of the Centre's Programme on Technology and Growth.

*I thank the British Academy, the ESRC, and the Andrew W. Mellon Foundation for financial support. Discussions with Louise Keely have aided my understanding of the issues below. Portions of this paper were presented at the LSE and the Cowles Foundation at Yale University, and as an Invited Lecture at the Indian and South-East Asian Meetings of the Econometric Society, Delhi, December 1998. Seminar and conference participants—in particular, participants in the UNU/WIDER project on Information Technology and Economic Development—have generously contributed insights.

Nontechnical Summary

That knowledge matters for economic growth is a proposition that is at once both old and new. But the claim now has its greatest relevance not in the growth impact of education, R&D, or science as traditionally construed, but in the increasing importance in national income of *knowledge-products*—computer software, new media, electronic databases and libraries, and Internet delivery of goods and services. I call these elements of the weightless economy knowledge-products, not necessarily because they are knowledge-intensive in production, but because their physical properties resemble those of knowledge. In this reckoning, the economy is knowledge-intensive, not just because of the quantity of knowledge used in production, but because of the quantity of knowledge-products consumed.

The Industrial Revolution of late 18th-century Britain exploited significant amounts of new scientific and engineering knowledge in steam engines, textiles manufacture, new methods of mining, and other production processes.

Between 1950 and 1990, the number of US scientists and engineers employed in research and development grew five-fold, from 200,000 to over 1 million.

These developments undoubtedly contributed to economies becoming more knowledge-intensive. They differ, however, from changes in the economy due to the increased importance of information technology, new media, libraries and databases, and genetic code intellectual property, among others. In the latter, the products consumed are weightless, i.e., they have themselves the properties of knowledge— infinite expansibility and irrelevance of physical distance; not so, by contrast, for the 18th-century Industrial Revolution or for the development of economies over the early 20th century.

Put succinctly, most notable about the very newest technologies is not just greater quantities at higher qualities of the same old stuff. Instead, it is that the economic properties of what is consumed differ importantly from those earlier.

This paper studies how that difference might matter for economic growth. In particular, it explores the notion that modern knowledge-product technologies reduce a relevant “distance” between consumers

and knowledge-based production. The conjecture is that, as a result, patterns of demand can have intricate and important effects on subsequent economic growth and technological developments.

The idea that modern technologies bring closer together consumers and knowledge production is distinct from the notion of general purpose technologies, which potentially affect the entire economic system at once, or that of skills-biased or directed technologies, which target particular segments of the worker distribution. All the latter focus on production-side characteristics. The distance reduction considered in this paper, by contrast, emphasizes the demand side.

(As just one example, recent information technology developments provide extreme cases where producers and consumers become ever closer. The **Linux** operating system is likely the clearest such instance—its current incarnation derives from improvements and refinements put in place by a large worldwide base of users. Consumer-side characteristics bear directly on the ongoing development of the technology. What this paper studies is not as stark or as direct, but will display similar features.)

By definition, the latest technologies have not been around for very long. Thus, convincing empirical time-series evidence on their impact will be difficult to obtain. Instead, this paper builds its argument and seeks its evidence on two distinct fronts: 1. real-world features of intellectual property protection in modern knowledge-products and 2. the failure of the putative Industrial Revolution in 14th-century China.

In the traditional view on knowledge and growth, because knowledge is costly to produce but easy to copy, societies require an institutional structure that protects the creators of knowledge and incentivizes them to create. A patents system is the standard such institution.

However, a substantial fraction of knowledge-products property is protected not by patents, but instead by copyright. Software and information compilations—significant in the IT value chain—have, until recently, largely been copyrighted, not patented. They are, therefore, legally the same as literary works. (However, in 1998 the US Patent and Trademark Office began to award patents for electronic-commerce business models—among them, Open Market’s Internet

marketing and payments system; CyberGold's attention-brokerage scheme; Priceline.com's buyer-driven, reverse auction model; Juno Online Services's advertisement-display techniques. Patent authorities reasoned that these were performed on computers, and thus were industrial, machine-driven processes.) Copyright protects an author's expression of an underlying set of ideas, but not the ideas themselves. Under most current legal systems, copyright is routinely awarded to any work showing originality, i.e., the work must not have been entirely copied, and must have had sufficient amount of labor, skill, and judgement involved in its creation.

Patents, in contrast with copyrights, require that an invention be novel, capable of industrial application, and innovative relative to the current state of the art; they are a much stronger form of intellectual property protection. For patents, the first to create the knowledge-product and thereby attain protection acquires the monopoly: There is a natural, winner-take-all feature in the dynamics.

Returning to copyright, all that is needed is to show that the work must have been the author's own creation. But this implies that there must be more than one way to implement an idea, as otherwise the work could not have been the author's own creation. An important economic implication follows from this: Copyright is not an intellectual property right allowing excludability and monopoly operation. Under copyright, others can freely copy portions of any work that are "critical", for which only one way exists to implement the idea—because those parts of a work could not have been the author's own creation.

In copyright, priority is also immaterial. Unlike with patents, it cannot matter whether a creator is the first to create the original instance of a copyrighted knowledge-product. Whatever central and critical idea is in the knowledge-product is unprotected, and thus can again be used by others.

Thus, just as knowledge-products are, in nature, easily copied, they are similarly so in a sense under copyright law. From the perspective of the traditional knowledge and growth models described above, copyright in IT should not work at all. Or, if copyright works, it does so in a way different from that routinely used in modelling knowledge protection in models of economic growth.

Consumers are not concerned in the same way rival producers are about being able to extract, for their own use, the critical good idea embedded in a copyright work left unprotected by law. (Although even this is contentious: decompiling a piece of software to see the critical idea—who decides when a particular subtlety is critical anyway—is almost certainly illegal.) It is the entire package of ideas and attributes that is more important. Consumers do not find as attractive a knowledge-product generated by alternative producers who bundle the central, unprotected attribute with other attributes differently implemented. Because those other attributes are protected by copyright, by definition they must be implementable differently. Consumers just want to enjoy the whole work conveniently, without inadvertently violating the law’s fair-use provisions by, say, viewing the work in a computer’s volatile random access memory. This act is impossible to perform without, in effect, making a copy—albeit temporarily—and thus infringing the rights of reproduction on a copyright work.

For knowledge-products protected by copyright, the distance between producers and consumers is smaller than when knowledge is protected by patent. The cost to a consumer of using a specific knowledge-product comes not only from the sticker price of the product, but from having to learn the norms, conventions, and subtleties involved in using the product.

Turn now to 14-th century China. Over the Sung (960–1126 C.E.) and Yuan (1127–1367 C.E.) dynasties, China became a technologically advanced economy. It had in place many of the same technical developments that later made possible the Industrial Revolution in late 18th-century Britain. Like Europe four centuries after, 14th-century China had solved the problems of making agriculture highly productive, efficiently manufacturing fine textiles, exploiting and applying kinetic and thermal energy, and producing high-quality materials for tools. Yet no Industrial Revolution occurred in China. An eminent economic historian calls this failure of 14-th century China to take off “the greatest enigma in the history of technology”.

The Chinese example illustrates a failure of technology-driven growth. In the standard view, technology is cumulative and path-dependent: Once technical advances occur, they should just keep

going. This, however, did not happen in Yuan China.

The puzzle lies not just in a comparison between China and Europe, but between a dynamic, innovative China before 1400 and a stagnant, regressing China after. Cultural explanations alone, therefore, cannot suffice.

China in the 14th century was not, of course, rife with software, genetic code products, and copyrights instead of patents. Instead, what it had was a severe mismatch between the supply and demand sides of technology. While before 1400 the state had exclusively subsidized the use of technologically-advanced goods, after that, it stopped doing so. Similar forces were at work in other parts of the world, but the extreme power and degree of exclusivity held in the hands of the state was unique to China. After 1400, that component of technology demand disappeared.

This paper builds on these observations to construct a model of knowledge-products and growth where, in equilibrium, the structure of demand determines patterns of aggregate economic growth. In some equilibria, depending on consumer attitudes, the knowledge-products industry languishes; in others, it is vital and fast-growing.

Although not central to this paper, implications on income distribution dynamics also follow from the model. Producers display Superstar outcomes in that their distributions of incomes are an emergent spreading-apart of the distributions of underlying attributes, and are related to the clusters or groupings in consumers.

The model is naturally interpreted as one of the software industry. Equally, however, it can be viewed to capture important elements of other idea-producing activities, such as the academic profession for one.

One policy implication from this study is the importance of training and education, not just in providing skills for work and production, but in providing a sufficiently strong demand base. Unless a demand side can be cultivated that appreciates and exploits sophisticated and advancing technology, economic growth can slow or, ultimately, fail to continue. Training and education need not be along narrowly-defined skills dimensions, but they can provide a double impact in strengthening both demand and supply sides of technical development.

Finally, patterns of demand emerge not just from developing skills and education in the population at large. They are affected also by government efforts at regulation. Government policy that strongly curtails private use of new technologies—excess taxation, insufficient access provision—can have adverse long-run consequences on growth in the economy.

The Weightless Economy in Economic Development

by

Danny Quah

LSE Economics Department

February 1999

ABSTRACT

Can the increasing significance of knowledge-products in national income—the growing weightless economy—influence economic development? Those technologies reduce “distance” between consumers and knowledge production. This paper analyzes a model embodying such a reduction. The model shows how demand-side attributes—consumer attitudes on complex goods; training, education, and skills for consumption (rather than production)—can importantly affect patterns of economic growth and development. Evidence from the failed Industrial Revolution in 14th-century China illustrates the empirical relevance of the analysis.

Keywords: China, copyright, growth, information technology, intellectual property, knowledge-product, patent, software, Superstars, tacit knowledge

JEL Classification: N15, O11, O14, O33

Communications to: D. Quah, LSE, Houghton Street, London WC2A 2AE.[Tel: +44 171 955-7535, Email: dquah@econ.lse.ac.uk]
[(URL) <http://econ.lse.ac.uk/~dquah/>]

1 INTRODUCTION

That knowledge matters for economic growth is an observation that is at once both old and new. But the claim now has its greatest relevance not in the growth impact of education, R&D, or science as traditionally construed, but in the increasing importance in national income of *knowledge-products*—computer software, new media, electronic databases and libraries, and Internet delivery of goods and services. These knowledge-products are so called, not necessarily because they are knowledge-intensive in production, but because their physical properties resemble those of knowledge.

This paper conjectures that knowledge-products rely importantly on demand for their ongoing growth and thus for their continued contribution to economic development. But because many of the relevant changes here are relatively recent, time-series evidence on their importance can be difficult to obtain.¹

Instead, the key idea the paper explores is that modern technologies associated with knowledge-products reduce the “distance” between consumers on the one hand and knowledge-based production on the other. As a result, demand-side factors—consumer attitudes on sophisticated goods; training, education, and skills for consumption (rather than production)—importantly influence patterns of technological development, and therefore growth and economic development more generally.

The idea that modern technologies bring closer together consumers and knowledge-based production is distinct from that of general purpose technologies (e.g., Helpman [12]), that potentially affect the entire economic system at once, or that of skills-biased or directed technologies (e.g., Acemoglu and Zilibotti [1]), that target particular segments of the worker distribution. All the latter focus on production-side characteristics; the distance reduction considered

¹ Elliott [10], Jorgenson and Stiroh [16], Jung, Oh, and Shin [17], Niininen [26], Sichel [36], and Wong [37] are useful recent studies in this regard.

in this paper, by contrast, emphasizes the demand side.²

To develop the argument, the paper considers two seemingly disparate sets of ideas:

1. 14th-century China was technologically advanced, and appeared to be on the verge of an Industrial Revolution—why did one not occur?
2. The properties of knowledge and information technology in general, but computer software in particular, make them unusual economic commodities—does the reduced distance between consumers and knowledge-producers suggest new models of knowledge and growth?

Section 2 briefly summarizes more traditional models of knowledge and growth, and describes how, in the framework of this paper, items 1.–2. are related. The section makes two points. First, 14th-century China provides an interesting example to study the relation between technology and growth. Advanced in ways that, according to some, paralleled the technological development of 18th-century England, China nevertheless failed to take off to sustained economic growth. Indeed, if anything, economic regress occurred over the subsequent four centuries. Second, the increasing importance in economies of information technology raises questions not easily dealt with in traditional models of knowledge and growth. In this paper’s analysis, these two—14th-century China and late 20th-century information technology—share a common feature: they emphasize the importance of consumer characteristics for determining the creation of knowledge-like goods and thus ongoing technological development.

Section 3 develops a model to analyze economic growth and development, emphasizing the importance of relating 1.–2. The model differs from others that deal with knowledge and growth in its focus

² Recent information technology developments provide extreme examples where producers and consumers become ever closer. The **Linux** operating system is likely the clearest such instance—its current incarnation derives from improvements and refinements put in place by a large worldwide base of users.

on how demand-side—rather than supply-side—considerations can be important for determining growth outcomes.³ In the model, demand-side factors produce a range of predictions interpretable as consistent with observations surrounding both 14th-century China and modern information technology. Consumer aversion towards using sophisticated technology can result in economic growth slowing or, in an extreme, failing to occur.

Section 4 concludes, summarizing the lessons from this study.

2 TECHNOLOGY, 14TH-CENTURY CHINA, AND GROWTH

Economists concerned with growth have long recognized the importance of technical progress.⁴ Leading theoretical models of technology and growth (e.g., Aghion and Howitt [2], Grossman and Helpman [11], and Romer [33]) have provided explicit economic descriptions of how technical progress, evolving through incentives and markets, drives economic growth.

To understand the insights from models of endogenous technology and growth, and to relate them to information technology (IT) in economic development, it helps to first situate the discussion in a broader literature.

2.1 Some economics of knowledge

Arrow [3] first formalized difficulties in the market for knowledge (see also the exposition in Dasgupta [5]). The property of knowledge most relevant for understanding IT is called *infinite expansibility* by David [7] (citing Thomas Jefferson [19]) and *nonrivalry* by Romer [33].

³ Of course, this statement should be read to refer only to a difference in emphasis. In any general equilibrium model, both demand and supply side contribute to the equilibrium outcome.

⁴ Keely and Quah [18] trace how this holds even for the neoclassical growth analyses of the 1950s. In this view, where many later writings differ is in their making explicit the economic incentives for generating technology, not necessarily in identifying new sources of growth.

These terms refer to the physical property of knowledge that it is not drawn down with use: an item of knowledge, once in place, can be used repeatedly, by many different users, in many different locations without the original piece of knowledge itself being degraded. So too computer software.⁵

Infinite expansibility implies knowledge has marginal cost equal to zero (for all practical purposes). But then, as long as marginal benefits remain positive, social efficiency requires that markets be flooded with free copies of that item of knowledge.

This last statement is normative. It goes beyond the predictive observation that knowledge, information, and ideas have an inherent tendency to disseminate freely and widely, i.e., is naturally *nonexcludable*. Instead, the statement stresses that their doing so is a necessary condition for outcomes to be (ex-post) efficient.

The difficulty, of course, is that should free dissemination occur, no economic agent would have the incentive to develop new pieces of knowledge in the first instance. An economic agent looking ahead sees the equilibrium price on the result of her efforts to be zero. Why not do something else then—appropriately compensating for one’s time—instead of producing new knowledge? Rational forward-looking agents will not produce knowledge, and thus the zero price for knowledge comes with zero quantity in equilibrium.

This market failure in knowledge reflects the tension between ex-ante incentive and ex-post efficiency. In public finance theory, such a public-goods problem might be solved by direct government intervention. For knowledge, the interventionist solution has observable counterparts in the institutions of *patronage* and *procurement*.

An example of patronage is public research funding, say by the ESRC in the UK or the NSF in the US. Grants are awarded based

⁵ The software example here and the intellectual property one below emphasize that knowledge and information in this analysis raise issues markedly different from those typically studied in, say, the “economics of information.” One can usefully analyze the economic impact of a word-processing package, process-controller software, genetic libraries, or database usage, without moral hazard, adverse selection, or contracts being at all central to the study.

on competitive proposals; research problems are determined by the investigator (not the source of the funding); and complete public disclosure of the research findings is agreed. Patronage is the dominant model of *open science* in Anglo-Saxon societies.

Examples of procurement, by contrast, are defense or space research, where government provides the resources to investigate questions that it, not the researcher, sets and where disclosure of the findings need not be total. Useful spinoffs—commercial radar, jet engine technology, the Internet—then enter the public realm of productive knowledge.

In describing patronage and procurement, I have followed the “3 P’s” terminology from David [8]. The third “P” is *property*, in the form of patents and copyrights, or intellectual property more broadly. Under property, society awards monopoly rights to the private creator of an item of knowledge. A market then forms in that the originator of the knowledge earns license fees from users; unlicensed copies of that protected knowledge are illegal.

All three P’s are second-best solutions to an endemic market failure in producing and distributing knowledge. The last of these, the intellectual property rights regime, by no means necessarily restores the first-best efficient outcome. Indeed, under different circumstances, one or the other of the three alternatives might be socially preferred (e.g., Wright [38]).

2.2 Knowledge in economic growth

How have models of knowledge-driven economic growth treated these theoretical difficulties in the production and distribution of knowledge?

Almost all formal analyses have sidestepped the more primitive issues described above, and assumed directly the existence of a system of intellectual property rights. In Romer [33], a stock of skilled workers—researchers, scientists, inventors, intellectuals—is available to generate ideas and new knowledge (Fig. 1). The greater the number of researchers and idea-producers, the faster the economy grows. Skilled workers face competing uses for their time and expertise: they can decide to be managers and lawyers, rather than research scien-

tists. In equilibrium, returns equalize across the different lines of activity for skilled workers, thereby determining the growth rate of the economy.

What is the payoff from producing knowledge? Upon obtaining a useful idea, the researcher obtains a patent for it, and sells that patent to a machine-maker. The machine-maker builds a machine around the idea. It then rents the machine to a monopolistically competitive manufacturing industry who—combining the machine with skilled managerial input—retail the final products to consumers. In equilibrium, the present discounted value of the stream of monopoly rent supports a positive equilibrium price on the patent.

Models in Aghion and Howitt [2], Grossman and Helpman [11], and others differ in focus and details from the description just given. For instance, Aghion and Howitt [2] study Schumpeterian dynamics, with the latest discoveries making obsolete earlier ones in a wave of creative destruction. Grossman and Helpman [11] consider schemes where economic goods proceed up a quality ladder, through the application of resources to R&D.⁶

In the framework just described, consumers come into contact with knowledge only through physical goods embodying the new technology. A chain of production runs from ideas being discovered through a patent protection scheme through a manufactured goods industry and finally to the consumer. Consumers are distanced from the producers of knowledge: indeed, consumers need never be aware of the infinite expansibility in knowledge underlying production.

Such a separation matters under two circumstances. First, with increasing pervasiveness of IT in general, and computer software in particular, the effective distance diminishes between consumers and knowledge production (Fig. 2). Consumer characteristics then have important impact on knowledge production and technical progress. A

⁶ There is sufficient commonality, however, across all such models that Jones [14] can persuasively criticize the entire class of “technology and growth” models by pointing out that while US scientists and engineers employed in R&D grew five-fold from 200,000 in 1950 to over 1 million in 1990, the growth rate of US GDP failed to increase by anything remotely comparable.

second circumstance is when severe enough mismatch appears somewhere along the chain between knowledge production and eventual consumption. China in the 14th-century, when an Industrial Revolution comparable to that in the West should have happened but did not, provides useful lessons here.

I now describe these two situations in greater detail, and use them to motivate the theoretical model in Section 3 below.

2.3 Knowledge-products and IT

The increasing importance of information technology shortens the effective distance between consumers and knowledge producers. By this, I don't mean that consumers directly eat the engineering and scientific knowledge being produced in R&D labs. Instead, I mean that software, although not in the usual sense scientific knowledge, has all the same physical and economic properties as traditional forms of knowledge. The same holds for information in electronic and biological libraries and databases.⁷ Thus, the resource-allocation difficulties for producing and distributing knowledge described above emerge with equal force for these *knowledge-products*. Consumer attitudes on using knowledge-products will increasingly affect the workings of whatever mechanism is in place for producing and distributing intellectual property.

A first complication in using standard models of knowledge and growth for analyzing IT is that knowledge-products are not typically made excludable by patent. Software and information compilations—significant in the IT value chain—have, until recently, largely been copyrighted, not patented. They are, therefore, legally the same as literary works.⁸ Copyright protects an author's expression of an

⁷ Quah [27, 30, 31] labels these elements of the *weightless economy*.

⁸ See, e.g., the UK Copyright, Designs, and Patents Act 1988. Section 1(2)(c) of the UK Patents Act 1977 explicitly excludes computer programs (Holyoak and Torremans [13]), although in 1998 the US Patent and Trademark Office began to award patents for electronic-commerce business models—among them, Open Market's Internet marketing and payments system; CyberGold's attention-brokerage

underlying set of ideas, but not the ideas themselves. Under most current legal systems, copyright is routinely awarded to any work showing originality, i.e., the work must not have been entirely copied, and must have had sufficient amount of labor, skill, and judgement involved in its creation.

Patents, in contrast with copyrights, require that an invention be novel, capable of industrial application, and innovative relative to the current state of the art. Patents comprise a much stronger form of intellectual property protection. For patents, the first to create the knowledge-product and thereby attain protection acquires the monopoly: There is a natural winner-take-all feature in the dynamics.

Returning to copyright, all that is needed is to show that the work must have been the author’s own creation. But this implies that there must be more than one way to implement an idea, as otherwise the work could not have been the author’s own creation. An important economic implication follows from this: Copyright is not an intellectual property right allowing excludability and monopoly operation. Under copyright, others can freely copy portions of any work that are “critical”, for which only one way exists to implement the idea—because those parts of a work could not have been the author’s own creation.

In copyright, priority is also immaterial. Unlike with patents, it cannot matter whether a creator is the first to create the original instance of a copyrighted knowledge-product. Whatever central and critical idea is in the knowledge-product is unprotected, and thus can again be used by others.

Thus, just as knowledge-products are, in nature, easily copied, they are similarly so in a sense, under copyright law. From the perspective of the traditional knowledge and growth models described above, copyright in IT should not work at all. Or, if copyright works, it does so in a way different from that routinely used in modelling knowledge protection in models of economic growth.

scheme; Priceline.com’s buyer-driven, reverse auction model; Juno Online Services’s advertisement-display techniques. Patent authorities reasoned that these were performed on computers, and thus were industrial, machine-driven processes.

Consumers are not concerned in the same way rival producers are about being able to extract, for their own use, the critical good idea embedded in a copyright work left unprotected by law.⁹ The entire package of ideas and attributes is the more important. Consumers do not find as attractive a knowledge-product generated by alternative producers who bundle the central, unprotected attribute with other attributes differently implemented. (Because those other attributes are protected by copyright, by definition they must be implementable differently.) Consumers just want to enjoy the whole work conveniently, without inadvertently violating the law’s fair-use provisions by, say, viewing the work in a computer’s volatile random access memory. (This act is impossible to perform without, in effect, making a copy—albeit temporarily—and thus infringing the rights of reproduction on a copyright work.)

For knowledge-products protected by copyright, the distance between producers and consumers is smaller than when knowledge is protected by patent. The cost to a consumer of using a specific knowledge-product comes not only from the sticker price of the product, but from having to learn the norms, conventions, and subtleties involved in using the product. This feature will be used critically in the model in Section 3 below.

Of course, knowledge-products are not all of IT. Computer hardware, in particular, can be understood within the standard models of technology and growth described above. But hardware in 1997 was only one-third of IT. In the foreseeable future it will likely become a progressively smaller fraction of the total.

Does this matter for poorer developing economies? In India, per capita annual income in 1995 was US\$340: The majority of the population of 900m lived on less than US\$1 a day. Yet, for the rest of the world, India hosts in Bangalore a major offshore high-tech software center. Indian software production in 1997 was a US\$2 billion industry, employing 260,000 people. The software industry in India has had revenues growing by 50% a year for the last five years, with

⁹ Although even this is contentious: decompiling a piece of software to see the critical idea—who decides when a particular subtlety is critical anyway—is almost certainly illegal.

over 60% generated as export earnings (in 1997 Europe accounted for 22% and the US 59% of Indian software exports; India's domestic software market generated revenues of US\$0.8 billion). Earnings of entry-level programmers in Bangalore might be less than 10% of their Western counterparts, but it is double that in the Indian Civil Service. This success has likely been due to a combination of factors: predominance of the English language in software programming and models; the ability to transmit value and work in the form of software code over the Internet; the timezone location of India relative to the west coast of the US, yet another major software center.¹⁰

An understanding of the ongoing role of IT in economic growth and development is therefore important. The discussion above suggest that it might be useful to consider theoretical extensions to existing models in two directions: first, to analyze the role of consumer attitudes towards sophisticated knowledge-products in determining aggregate patterns of growth and development; second, to study the aggregate growth effects of alternative systems of intellectual property rights or, more accurately, systems for producing and distributing knowledge and knowledge-products. In this paper, I take up only the first, and leave the second for subsequent research.

2.4 China in the Sung and Yuan dynasties

China from the 14th through 19th centuries provides an interesting case study for technology-driven economic growth.

Over the Sung (960–1126 C.E.) and Yuan (1127–1367 C.E.) dynasties, China became a technologically advanced economy. It had in place many of the same technical developments that later made possible the Industrial Revolution in late 18th-century Britain. Like Europe four centuries after, 14th-century China had solved the prob-

¹⁰ These “death of distance” factors are not the only ones at work, however. Many Indian software professionals have also begun working in the US under the new high-tech worker visas agreement, lobbied for by many Silicon Valley businesses. The number of visas permitted has increased from 65,000 in 1998 up to 115,000 by 2001, after which the limit is scheduled to drop back to 65,000.

lems of making agriculture highly productive, efficiently manufacturing fine textiles, exploiting and applying kinetic and thermal energy, and producing high-quality materials for tools. Yet no Industrial Revolution occurred in China.¹¹ Why?

In China, blast furnaces for casting iron and refining wrought from pig iron were in use by 200 B.C.E.; Europe did not achieve this until the late 14th century. Blast furnaces were being fuelled by coke from the early 11th century in China for smelting iron, it was only in the 18th century that similar technologies became available in Europe. Indeed, by 1078 China was producing iron per capita at rates comparable to Europe in 1700; total Chinese iron output in 1078 exceeded that for Britain in 1788. In China the price of iron relative to grain fell from 6.3 in 997 to 1.8 in 1100; iron did not become as cheap in England until 1700. The iron plow was introduced in the 6th century and was adapted for wet-field rice cultivation by the 9th century.

China’s water power development paralleled that in Europe. By 1280 the use of the vertical water wheel was widespread. Gunpowder had been invented in China before the 11th century, and a thriving chemicals industry existed to provide, among other things, powerful explosives for military and industrial application. Europeans did not have gunpowder until the late 13th or early 14th century, but their subsequent developments of it—for instance, corning the powder rather than retaining its use in fine-grain form as did the Chinese—soon gave them the military edge.

Spinning wheels for textiles appeared in China in the 13th century, about simultaneous with the West, but advanced faster than in Europe from adapting earlier Han dynasty expertise in fine silk weaving. Water-powered spinning machines by 1100 were as advanced as those in early 18th-century Europe, but it was not until the British

¹¹ The descriptions I give below are documented and discussed further in Derry and Williams [9], Jones [15], Landes [22], McNeill [24, Ch. 2], and Mokyr [25]. After describing Sung Chinese technical advances, Mokyr [25, p. 213] concludes that “China was about ready to undergo a process eerily similar to the great British Industrial Revolution”.

Industrial Revolution that Europe enjoyed similarly systematic exploitation of central power sources for textile working.

Chinese water clocks of the 11th century were, according to some, more accurate than contemporary European timekeeping devices.¹² Su Sung’s water clock of 1086 was 40 ft. high and showed not just the time but a range of astronomical indicators, including the positions of the moon and planets: it is widely viewed by historians as a technological masterpiece.

After its invention in China from before the first century, paper took over 1000 years to reach the West (and then only after some Chinese paper-makers were taken prisoner by Arabs in a battle for Samarkand). China was using block-printing by the late 7th century and porcelain moveable type from 1045, four centuries before Gutenberg. Metal type was available from the early 13th century. The invention of the compass around 960 allowed Chinese junks to be the most advanced navigators in the world. On top of this, Chinese ships were from 1400 constructed using watertight bouyancy chambers, a superior technique of ship construction not adopted by the West until the 19th century.

Yet, all this—four hundred years earlier than the West—sparked no takeoff remotely comparable to the late 18th century Industrial Revolution in Britain and Europe. Indeed, by the 1800s, China was widely recognized to be technologically backward relative to England. The Chinese were aware of this from as early as 1600. The 1842 Opium War between Britain and China showed the power of superior technology in exacting a punishing defeat on a once-proud empire. Technology did not simply stagnate; absolute regress occurred. By the 16th century, Chinese time measurement had become again primitive, despite Su Sung’s waterclock of five centuries earlier.

A number of puzzles emerge: Why did technological innovation wither away towards the end of the Yuan dynasty? If technology is cumulative and path dependent—with success building upon success—why did China not continue to be the world’s technology leader?

¹² Landes [22] points out that European mechanical timepieces soon showed their superiority. Water clocks were not durable; their mechanisms wore down and became unreliable faster.

And, after themselves acknowledging their backwardness from the 17th century on, why did the Chinese not seek to copy the superior Western technology they could readily observe?¹³

The puzzle lies not just in a comparison between China and Europe, but between a dynamic, innovative China before 1400 and a stagnant, regressing China after. Cultural explanations alone, therefore, cannot suffice.

The hypothesis explored in this paper is a re-interpretation of one advanced by Mokyr [25, Ch. 9]: a mismatch appeared between patterns of supply and demand for technology. In contrast with the situation in Europe, the Chinese state was a significant player for determining demand for ideas and for technology. When the Chinese government withdrew support, private enterprise was unable to step in sufficiently quickly. By contrast, no one government controlled all of Europe, and even within nations, the state was but one of a number of competing customers for ideas and technology.¹⁴

The Chinese state and bureaucracy changed after 1400. For one,

¹³ Mokyr [25, p. 209] calls the failure of 14th-century China to maintain its position of superiority “the greatest enigma in the history of technology”.

¹⁴ Krussel and Rios-Rull [21] give an alternative formalization, also citing Mokyr [25]. In their work, the state and bureaucracy—vested interests—are averse to allowing new technology to disseminate. Overly restrictive control can then lead to stagnation. In the development below, the state is viewed as just one among others on the demand side for technology (although admittedly not very many others). Europe too had reactionary governments and ruling classes with vested interests. Unlike China, however, no single state controlled all of the European continent. Europe managed to prosper and eventually underwent an Industrial Revolution while China did not. Thus, it is not vested interests alone that were critical, but instead the overall pattern of demand. I interpret the Chinese situation after 1400 as a change “of preferences, of attitudes toward technological change and its consequences”. Mokyr [25, p. 232] used this phrasing to compare Europe and China after 1400, but it seems to me to apply equally well to several other explanations he provides.

when the Mongols conquered China and established the Yuan dynasty, the switch to the non-iron based Mongol military technology eventually reduced iron demand to a degree significantly affecting related products. The coal industry, which had been active from the early 11th century, fell into decline. Mines came to exploit only shallow excavations; operators discarded the use of machinery for pumping, ventilation, or moving out the mined product. Mining changed from a state of relatively advanced technology to one relatively primitive.

The government in China viewed technology and knowledge as its exclusive domain. Time itself belonged to the Emperor, and the government attempted to monopolize both its measurement and the use of the calendar. It was the Emperor who instructed government officials to commission and use Su Sung's waterclock masterpiece. Before the Mongols, the Imperial government established state-owned iron foundries to promote using iron implements. The Sung government encouraged farmers to use improved technology by directly providing financial incentives: from the Han period onwards, the state provided peasants with the physical capital necessary for technological improvement, including tools and draft animals, and promoted the use of more advanced plows. Early on, the Yuan dynasty continued to actively promote the use of new textiles.

Therefore, before 1400, the state in China played an important role in generating demand for technology. When this Imperial support was withdrawn after 1400, technology stagnated and regressed.

In Europe, by contrast, technical change arose mostly from private and explicitly commercial demand. Governments and rulers were typically secondary and passive. James Watt and Matthew Boulton had immediate customers for their steam engines in Cornish tin-mines and a fast-growing iron industry. For them it was the supply side—skilled labor, specialized material and tools; financial support—that provided the binding constraints.

The example of 14th-century China shows that the demand for knowledge and technology matters for growth and development. In China the supply of technology was present, but changing patterns of demand led to technology first growing rapidly, then languishing, and finally actually regressing. Economic growth failed as a result.

Viewing technology in a way that distances the consumers and producers of knowledge hides this connection between demand and supply. The characteristic with information technology now is that, once again, the consumer side comes into direct contact with technology and knowledge-products. On the supply side of IT, technology has, arguably, advanced beyond consumers' comfort threshold; the demand side is suspicious of IT because of the implied potential loss of privacy, its complexity of use, and the perceived fragility of the technology. How growth and development will unfold depends on the response of the demand side in using sophisticated goods, the intellectual property environment that fosters ongoing development in technology, and the equilibrium reactions of the supply side.

To analyze these interactions, we turn now to a model that formalizes the idea that the demand side matters for technological progress and in determining equilibrium growth outcomes. The model is a simplified, steady-state version of that studied in Quah [32].

3 THE MODEL

Assume an economy having distinct populations of producers and consumers, and proceeding in discrete time $t = 0, 1, 2, \dots, \infty$.

Producers will maximize expected present discounted value of revenues over occupational choice—working as an entrepreneur in the knowledge(-products) industry or at a safer job. Knowledge-products—software or ideas—can have either established high quality or be experimental, with the possibility of emerging as success or failure. The payoff to working in knowledge-products is determined endogenously in equilibrium.

Consumers choose which knowledge-products to purchase. Consumers are characterized by heterogeneity in their aversion towards using the knowledge-product. Depending on the distribution of this hedonic cost, a range of equilibria can emerge. Put differently, consumer attitudes towards the use of sophisticated technologies determines the growth outcome in equilibrium.

In some equilibria, the knowledge-products industry languishes; in others, it is vital and fast-growing.

Although not central to this paper, implications on income dis-

tribution dynamics also follow from the model. Producers display Superstar outcomes in the sense of Rosen [34] and MacDonald [23]: income distributions are an emergent spreading-apart of those in underlying attributes, and are related to the clusters or groupings in consumers.¹⁵

The model is naturally interpreted as one of the software industry. Equally, however, it can be viewed to capture important elements of other idea-producing activities, such as the academic profession for one.

3.1 Producers

Producers are drawn from two-period lived overlapping generations that are born ex-ante identical. Every generation contains P potential producers, each seeking to maximize over choice of occupation and output supplied the present discounted value of expected profits, with discount factor $\beta \in (0, 1)$.

Fig. 3 shows the demography and timing of decisions. When members of generation t are born at the beginning of period t , they decide whether to become regular workers or to develop ideas for the knowledge or θ industry. If they choose to be regular workers, they earn the outside option $w(t)$ at time t . If, however, a member of generation t becomes an idea-developer at time t , she produces a knowledge-product with quality $\theta_x(t) \in (0, \infty)$. (Ideas and knowledge-products are synonymous in the model.) The x subscript denotes *experimental*, and θ_x is the same across all active experimenting idea-developers— young and old—in a given time period. Quantity θ_x evolves across time periods; its dynamics are described below.

When old, those in generation t who were previously regular workers can either continue to be regular workers—now at the outside option wage $w(t + 1)$ —or they can start to be idea-developers. If the latter, because they are new at idea-developing, the old of generation t produce $\theta_x(t + 1)$, as do the young of generation $t + 1$.

Those in generation t who when young were in the knowledge-

¹⁵ The dynamics resemble those for income distributions across countries (Quah [28, 29]).

products business are seasoned developers by time $t + 1$. If such an agent continues to develop knowledge-products when old, she produces an idea that is a random variable realizing as either $\theta_s(t + 1)$ (success) with probability π or $\theta_f(t + 1)$ (failure) with complementary probability $1 - \pi$.

Success and failure outcomes are distinguished by

$$\forall t : \quad \theta_f(t + 1) < \theta_s(t + 1),$$

and are distributed independently across seasoned idea-producers. The (θ_s, θ_f) pair can be attained only by those who have undergone the experimental stage, i.e., an initial release is necessary before a knowledge-product matures to become a success or a failure.¹⁶

Extensions of this work would usefully consider replacing this exogenous $(\pi, 1 - \pi)$ probability mechanism for describing successes and failures. Entrepreneur talent and skills will, in general, determine success, although other factors might matter as well. For instance, a low-technology installed base or—more interesting for the perspective developed in this work—consumption network externalities could lead to economic successes, θ_s , that are not identical with technological ones (Arthur [4] and David [6]). The current work assumes away those effects, but planned future research will seek to integrate them with those emphasized here.

At t the experimental knowledge-product has quality $\theta_x(t)$ borrowing from the current state of acknowledged successes $\theta_s(t)$, i.e.,

$$\theta_x(t) = \lambda_x \theta_s(t), \quad \text{with } 0 < \lambda_x < 1. \quad (1)$$

Let $P_x(t) \in [0, P]$ denote the number of young in period t working as idea-producers. I will assume that when P_x differs from zero, both it and P are sufficiently large to allow a law of large numbers across producers.

Assume that success quality evolves as:

$$\forall t \geq 0 : \quad \theta_s(t + 1) = \begin{cases} \lambda_0 \theta_s(t) & \text{if } P_x(t) = 0, \\ \lambda_s \theta_s(t) & \text{otherwise,} \end{cases} \quad (2)$$

¹⁶ We might think of this in the vernacular as requiring a beta version or a version 1.0 before a knowledge-product becomes established.

with $\lambda_s \geq 1 > \lambda_0 > 0$ and $\theta_s(0) > 0$ given. Equation (2) specifies, as a simplification, that growth in θ_s does not vary smoothly in P_x , but depends only on whether *some* young are working as idea-producers. When those exist, growth is at rate λ_s ; otherwise, growth is lower at rate λ_0 . Indeed, since $\lambda_0 < 1$, the economy actually regresses when there are no experimenting idea-producers.

A different and perhaps more natural specification might have growth rates continue to increase as P_x rises. Equation (2), appropriately monotone increasing in P_x , simplifies the analysis without giving up that intuition. The constant proportionality in (1) means that with a lag the growth rate in experimental quality equals that in success quality, regardless of which branch of (2) is active.

Because $\lambda_0 < 1$ technical regress occurs in the absence of experimental activity. This, however, is inessential: what matters is only that $\lambda_s > \lambda_0$ so that we could have, for instance, $\lambda_0 = 1$. Then, no technical regress occurs, even when idea-producers are inactive for several periods: Progress simply picks up where it had earlier left off.

Assume that experimental idea-products have quality intermediate between that of failures and successes, $\theta_x(t) \in (\theta_f(t), \theta_s(t))$, and that there is free entry into the knowledge-product industry.

We have discussed knowledge-product quality. Turn now to the costs of producing knowledge-products in quantity. From Arrow [3], and following Arthur [4], Krugman [20], and Romer [33], it is natural to suppose that such marginal costs are either zero or falling. However, to focus on the effects that are novel here, assume not the standard increasing returns, but that it costs an idea-producer $\phi_t(G)$ at time t to supply G units of her knowledge-product, with $\phi_t(0) = 0$, $\phi_t' > 0$, and $\phi_t'' > 0$. Costs ϕ_t are, therefore, taken to be identical across θ , although varying through time. Invariance in θ of marginal costs then gives increasing returns in the θ direction, although there is no increasing returns in the quantity G direction.

At any time t three kinds of knowledge-products, $\theta_f(t)$, $\theta_x(t)$, and $\theta_s(t)$, are potentially traded. Denote their corresponding spot prices $p_f(t)$, $p_x(t)$, and $p_s(t)$. This notation imposes that knowledge-products in the same class command the same price. In a given time

period an idea-producer with knowledge-product θ receives profit

$$R = \max_{G \geq 0} pG - \phi(G), \quad (3)$$

so that optimal G and R increase in p . Indeed, from the envelope theorem, the rate at which R increases in p is $(\phi')^{-1}(p)$ and is itself increasing in p . The more convex is ϕ , the more skewed is the implied distribution of revenues across producers for a given configuration of prices. Because $\phi(0) = 0$, we must have R nonnegative. Label R_f , R_x , and R_s the profits corresponding to the different classes of knowledge-products, and let the associated supply decisions be G_f , G_x , and G_s . If $p_f < p_x < p_s$, then so too $R_f < R_x < R_s$ with absolute differences magnified from those in p .

Finally, assume that the outside option $w(t)$ and the cost structure ϕ_t evolve exogenously. For existence of equilibrium, it will be convenient to assume that w in each period falls in an appropriate intermediate range. The conditions determining that range don't have a directly interesting economic interpretation; so I assume here only that the outside option w is neither too large nor too small.

To summarize the essential elements above, write the producer's problem as the value equation:

$$\begin{aligned} R_x + \beta [\pi \max\{R'_s, w'\} + (1 - \pi) \max\{R'_f, R'_x, w'\}] \\ = w + \beta \max\{R'_x, w'\}, \quad (4) \end{aligned}$$

where $'$ denotes values in the second period of life. When (4) is satisfied, producers are indifferent across occupations. The demand side then determines P_x .

3.2 Consumers

Consumers live for one period, at the beginning of which each receives exogenous income Y . The number of consumers is C constant through time. Heterogeneity across consumers is indexed by $\nu \in (-\infty, \infty)$ distributed following cdf \mathbb{F} . Population heterogeneity and incomes evolve exogenously so that at t they are \mathbb{F}_t and $Y(t)$ respectively.

Two kinds of goods can be consumed: knowledge-products and a numeraire composite commodity. Consumption of the first is time-consuming, so that—as in Shaked and Sutton [35]—in a lifetime at most one knowledge-product from a single class can be consumed.

Consumer type ν solves the problem:

$$\begin{aligned} \max_{\substack{c \geq 0 \\ \theta \in \{\theta_f, \theta_x, \theta_s\} \cup \emptyset}} & U(c, \theta \mathbf{1}_\theta) \\ \text{s.t. } & c + (p_\theta + \nu) \mathbf{1}_\theta \leq Y \end{aligned} \tag{5}$$

with

$$U_c > 0, U_\theta > 0 \quad \text{and} \quad \lim_{c \rightarrow 0} U(c, \theta) = 0 \text{ for fixed } \theta, \tag{6}$$

and $\mathbf{1}_\theta$ denoting the indicator function on $\{\theta_f, \theta_x, \theta_s\}$.

The budget constraint in (5) says that consuming θ entails payment on top of the sticker price p_θ the further cost ν . Condition (6) says that consumers value both composite commodity quantity and knowledge-product quality θ , but leaves unrestricted the substitution propensities across the two kinds of goods.

The higher is type ν , the more costly is consuming the knowledge-product. Although ν differs across consumers, it is invariant to quality θ . This has two interpretations. First, following Rosen [34], ν can be viewed as measuring the opportunity cost of time. The act of consuming knowledge-products—like enjoying opera or athletic performances—takes time. The opportunity cost of that time does not depend on whether the opera or knowledge-product is high- or low-quality, only on whether or not the consumer attends the performance. High- ν consumers would then, other things equal, have a further reason for selecting high-quality knowledge-products.

Second, ν can be viewed as parameterizing the cost of learning. The higher is ν , the less easily the consumer learns to access and appreciate a particular knowledge-product.¹⁷ Put differently, ν describes *tacit knowledge*. It is knowledge that is specific to

¹⁷ For knowledge-products that are computer software, there might be a specific set of conventions—keyboard and mouse actions, menu configurations, and so on—one has to learn that is common to all

individuals—hence its distribution across the population—and not infinitely expansible. Tacit knowledge is what allows users to access and exploit knowledge-products. In this second interpretation, negative ν can be interpreted as a subsidy to learning.

In either interpretation, ν invariant over θ means that consumers find it no easier or harder to appreciate an experimental θ_x than they do an established success θ_s . Thus, in the model the inherent quality of the knowledge-product is fully described by θ in the utility function, not by its potentially affecting the budget constraint in (5). Type ν describes consumer attitudes towards knowledge-products altogether, not towards any one of them.

Consumer type ν can always choose to consume no knowledge-product at all, in which case ν is irrelevant and all income is spent on the composite commodity c . Otherwise, composite commodity consumption $c = Y - p_\theta - \nu$, whereupon utility is $U(Y - p_\theta - \nu, \theta)$.

Call C_x the number of consumers choosing θ_x . Similarly, let C_f and C_s be the number of consumers choosing θ_f and θ_s respectively.

3.3 Equilibrium

Producers and consumers understand the laws of motion of the different variables in the economy, and in particular those (unspecified in the discussion thus far) in the exogenous quantities w , ϕ , \mathbb{F} , and Y .

An equilibrium is a sequence of spot prices and occupation and supply decisions such that all consumers maximize utility and all members of the producer populations maximize expected present discounted value of profits. Formally, we seek sequences

$$\{p_f(t), p_x(t), p_s(t), P_x(t), G_f(t), G_x(t), G_s(t) : t = 0, 1, \dots, T\}$$

satisfying (3), (4), and (5), given the exogenous evolution of outside

software, but the quality of screen presentations and metaphor developments in the product differ across θ 's. For knowledge-products that are ideas, one might have to learn the technical language in which such ideas are expressed, but that language is invariant to the quality of the ideas themselves.

options $w(t)$, cost structure ϕ_t , consumer attitudes \mathbb{F}_t , and incomes $Y(t)$.

The discussion is considerably simplified if we focus on equilibria bearing a stationarity property. It is convenient for this to take $\lambda_s = 1$, i.e., the economy even with ongoing successful innovation shows zero growth (without ongoing innovation, on the other hand, the economy decays at rate λ_0). Define a *stationary equilibrium* to be an equilibrium where for all t

$$p_f(t+1) = p_f(t), \quad p_x(t+1) = p_x(t), \quad p_s(t+1) = p_s(t),$$

$$\mathbf{P}_x(t+1) = \mathbf{P}_x(t),$$

and

$$G_f(t+1) = G_f(t), \quad G_x(t+1) = G_x(t), \quad G_s(t+1) = G_s(t).$$

Assume that the exogenous quantities are, similarly, time-invariant, i.e.,

$$w(t+1) = w(t), \quad \phi_{t+1} = \phi_t, \quad \mathbb{F}_{t+1} = \mathbb{F}_t, \quad Y(t+1) = Y(t).$$

Hereafter, we study only such stationary equilibria.

Free entry implies that in each period $R_x \leq w$. Moreover, since $\theta_f < \theta_x < \theta_s$ and utility is increasing in θ , we must have $p_f < p_x < p_s$, and therefore $R_f < R_x < R_s$. The producer's value equation (4) then becomes

$$R_x + \beta\pi \max\{R'_s, w'\} = w + \beta\pi w' = (1 + \beta\pi)w. \quad (7)$$

However, if $R'_s \leq w'$ then $R_x = w \geq R_s$, which is a contradiction. Thus, in each period $R_f < R_x < w < R_s$. New idea-developers accept low earnings temporarily so that they can later potentially become high-earning Superstar successes. Those idea-developers whose experiments turned out to be failures exit the market and work the outside option.

In equilibrium, therefore, no θ_f 's are produced. Call ζ_s the state of consuming θ_s while paying price p_s ; similarly define ζ_x . The consumer's problem (5) reduces to:

$$\max_{\theta \in \{\theta_x, \theta_s\} \cup \emptyset} \{U(Y - p_s - \nu, \theta_s), U(Y - p_x - \nu, \theta_x), U(Y, 0)\}. \quad (8)$$

For fixed Y and (p, θ) , the function $U(Y - p - \nu, \theta)$ decreases with ν . Taking variation in type ν , how costly it is to use knowledge-products, the three functions $U(Y - p_s - \nu, \theta_s)$, $U(Y - p_x - \nu, \theta_x)$, and $U(Y, 0)$ can be graphed as in Figs. 4–6. Starred and unstarred schedules in Figs. 4–6 represent either s or x choices respectively: whether the s schedule, say, is higher or lower than the x schedule at low values of ν turns out to depend on characteristics of the utility function and the range of variation in θ —see also the discussion surrounding Figs. 7–8 below. In all cases, the solution to (8) traces out the upper envelope of the three schedules in U .

Figs. 4–6 show that exactly three outcomes are possible: the graphs of $U(Y - p_s - \nu, \theta_s)$ and $U(Y - p_x - \nu, \theta_x)$ might intersect, or they might not (Fig. 4). If the former, the intersection might occur below $U(Y, 0)$ (Fig. 5) or above (Fig. 6). (In principle, multiple intersections might occur. Conditions on U that rule out such multiplicity are available, but don't seem to add much insight. Thus I simply take the single intersection possibility as primitive.)

Taking (Y, p_s, p_x) as given define ν_{int} to be the intersection in ν of $U(Y - p_s - \nu, \theta_s)$ and $U(Y - p_x - \nu, \theta_x)$, i.e.,

$$\nu_{int} \stackrel{\text{def}}{=} \{\nu \ni U(Y - p_s - \nu, \theta_s) = U(Y - p_x - \nu, \theta_x)\}.$$

Let ν_{max} be the larger of the intersections of $U(Y - p_s - \nu, \theta_s)$ and $U(Y - p_x - \nu, \theta_x)$, respectively, with $U(Y, 0)$, i.e.,

$$\nu_{max} \stackrel{\text{def}}{=} \max(\{\nu \ni U(Y - p_x - \nu, \theta_x) = U(Y, 0)\}, \{\nu \ni U(Y - p_s - \nu, \theta_s) = U(Y, 0)\}).$$

From (6), threshold ν_{max} always exists, even if ν_{int} might not.

The cases depicted in Fig. 4 and Fig. 5 are, in effect, the same. In either Figure, only two types of consumption behavior occur: one or the other of θ_x and θ_s (but not both), or no knowledge-product at all. However, neither situation in Fig. 4 and Fig. 5 is sustainable. To see this, note that no idea-producer works at θ_x for $R_x < w$ when in the next period θ_s experiences zero demand. Conversely, no idea-producer achieves the maturity to produce θ_s if no demand is expressed for θ_x when she is young. By contrast, Fig. 6 describes sustained technological progress.

China in the 14th century is, in this analysis, an example of Fig. 4 and Fig. 5. From a position where technical advances regularly occurred and were sustained (Fig. 6), the ν distribution shifted upwards, leading to where the induced demand for technical advances became progressively narrower and finally non-existent.

Turn to Fig. 6, where technical advances can be sustained. Here, again, there are two cases: $U(Y - p_s - \nu, \theta_s)$ can intersect $U(Y - p_x - \nu, \theta_x)$ from above (Fig. 7) or from below (Fig. 8). Which of the two obtains depends on the interaction between possibilities of substitution across c and θ and on the configuration of quality θ and equilibrium price p . (For instance, CES functional forms for U can give either Fig. 7 or Fig. 8, depending on whether the elasticity of substitution exceeds 1. The standard Cobb-Douglas special case gives Fig. 7.)

I have labelled Fig. 7 the *Learning Society* because in it the middle class—as measured by their type ν falling in a range intermediate between the very low and very high—choose to consume the experimental θ_x . It is this middle group of consumers, whose costs of using the knowledge-product are intermediate, that supports innovation and experimentation in the knowledge-products industry.

By contrast in the *Conservative Society* of Fig. 8 the middle class chooses only established successes. It is now those with low ν 's—those who find it least costly to use knowledge-products—who demand experimental knowledge-products.

The distinction between the two societies in Figs. 7–8 is substantive. If, for example, governments seek to subsidize learning by reducing consumers' opportunity costs (induce negative ν), they will end up only subsidizing consumption of established successes in Fig. 8, not of the experimental knowledge-products, as would happen in Fig. 7. (Of course, increasing the demand for θ_s does have, in general, positive knock-on effects for θ_x since dynamic rewards then increase for those in the knowledge-products industry.)

To see that a stationary equilibrium exists for both the Learning and Conservative Societies, we proceed in steps. In either society, no θ_f is produced and we can set $C_f = 0$. In the Learning Society, demand for successful knowledge-products is $C_s = F(\nu_{int})$ while supply is $\pi P_x G(p_s)$; demand for experimental knowledge-products is $C_x =$

$F(\nu_{max}) - F(\nu_{int})$ while supply is $P_x G(p_x)$. Holding fixed the number of experimenting idea-developers P_x , an increase in p_s reduces $F(\nu_{int})$ and increases $G(p_s)$; an increase in p_x reduces $F(\nu_{max}) - F(\nu_{int})$ and increases $G(p_x)$. Thus, for fixed P_x , a price pair (p_x, p_s) clearing the knowledge-product markets always exists. The same conclusion holds for the Conservative Society, remembering to reverse both the interpretations and the supplies that correspond to the different demands $F(\nu_{int})$ and $F(\nu_{max}) - F(\nu_{int})$.

Of course, for arbitrary P_x the (p_x, p_s) pair that clears knowledge-product markets need not satisfy the producer value equation (7). But the left-hand side of (7) varies continuously in P_x so that provided w falls in an appropriate intermediate range (as discussed above), then there always exists an equilibrium sub-population P_x of idea-producers in $(0, P)$ such that the knowledge-product market-clearing prices then also imply revenues (R_x, R_s) satisfying (7).

In the model, the exact value of the outside option w , once it falls within the appropriate range, does not matter for the equilibrium. This feature is an artifact, however, of the discreteness in the dynamics in (2). More generally, the relation between w and the rate of growth will behave as in Romer [33], where w is endogenously determined by an equilibrium condition depending on the productivity of the alternative (not knowledge-producing) sector.

Summarizing, a stationary equilibrium always exists and has the properties described in Figs. 7–8.

3.4 Income mobility and inequality: Distribution dynamics

Equilibrium, regardless of whether it is in the Learning or Conservative Societies, displays certain properties common to all Superstars models. Other features, however, are a little surprising.

The central Superstars result holds in the model. From consumers' optimization, price p is increasing in quality θ . Since production costs ϕ are invariant to θ , high-quality idea-producers optimally produce more than low-quality ones. The reasoning surrounding (3) then gives that profits are a convex function of quality, so that there is the usual skewed income distribution, already familiar from MacDonald [23] and Rosen [34].

Only a few Superstars (πP_x) survive to earn very high rewards. At the same time, many more experimenters (P_x) enter the industry at low immediate earnings, anticipating that they too might become Superstar producers in the future. Many of them fail, however, and exit the industry.

Finally, since $R_x < w$, no producer begins a career as an idea-developer in midlife. A regular worker when young remains a regular worker throughout.

On the consumer side, as expected from the standard Superstars reasoning, those with high opportunity or learning costs ν —on the extreme right of Figs. 7–8—spend none of their resources on the knowledge-product. But then, however, unlike in MacDonald [23] and Rosen [34], there is no necessary monotonicity in consumption patterns.

To be clear on this, recall the discussion of opera consumption in MacDonald [23] and Rosen [34]. High ν consumers find their time opportunity costs too high and consume no opera at all. Slightly lower ν consumers never find it worthwhile to try out experimental (potentially) low-quality opera, and instead always consume just the established successes. Finally, the least discriminating consumers, with lowest ν 's, only go to cheaper experimental, low-quality opera.

Here, by contrast, intermediate consumers with $\nu \in (\nu_{int}, \nu_{max})$ do not always consume just the high-quality knowledge-product. In the Learning Society of Fig. 7, those middle- ν consumers experiment—the equilibrium price p_x turns out to be sufficiently low to induce them to do so. It is only in the Conservative Society of Fig. 8 that the middle- ν consumers choose to consume high-quality established successes θ_s .

Although somewhat outside the model, a leapfrogging interpretation is also available. In the Learning Society Fig. 7, those with smallest ν 's—who have the lowest opportunity cost of time—bypass the lower-quality experimental products and latch right on to the high-quality ones. If we identified such consumers as the least-developed economies (ignoring the assumption that, in the model, Y is the same across all consumers), Fig. 7 says that they immediately jump to the knowledge frontier. At the same time, however, those economies already relatively developed, with ν 's in an intermediate range, choose

to learn and use only experimental technologies (because the price on those is relatively low in equilibrium). In the Conservative Society Fig. 8, the situation is reversed. The least-developed economies only use lower-quality experimental knowledge-products, while the relatively more-developed ones, only the established high-quality successes.

4 CONCLUSION

This paper studied the importance of the demand side in technology and growth, with special reference to the development of information technology.

Knowledge and technology development have long been recognized to be important in economic growth. This paper considered the conjecture that what matters most in modern technologies—information technology and other knowledge-like goods—is not that the resulting goods continue to be knowledge-intensive, but that they have many of the same physical and economic properties as knowledge itself.

Because these changes are relatively recent, however, time-series evidence on their importance is difficult to obtain. By the same token, it is hard to disentangle empirically what factors and institutional structures will matter for their ongoing evolution.

This paper took instead the view that weightless-economy changes are significant because they shorten the relevant “distance” between consumers and producers of knowledge-products. This then allows analysis on two fronts. First, what is special about the knowledge embodied in IT for economic growth? What interesting economic features do traditional models of knowledge and growth miss? For one, in reality, copyrights matter for IT more than do patents for protecting knowledge as intellectual property. But copyrights protect in ways fundamentally different from patents; compared to the latter, the level of protection afforded is weak. One reason this can be sustained is that the cost to consumers of using IT is not just the sticker price of the product, but the associated costs of learning to consume the knowledge-product.

Second, what does empirical evidence show in *other* situations

where similar distance-reduction has occurred? Clearly, 14th-century China did not have computers and information technology. But, as argued above, the tensions manifest between mismatched demand and supply sides of technology can provide useful lessons.

Building on these observations, the paper then developed an analytical model to show how demand side factors can produce a range of predictions consistent with both 14th-century China and the current information technology industry.

One policy implication from this study is the importance of training and education, not just in providing skills for work and production, but in providing a sufficiently strong demand base. Unless a demand side can be cultivated that appreciates and exploits sophisticated and advancing technology, economic growth can slow or, ultimately, fail to continue. Training and education need not be along narrowly-defined skills dimensions, but they can provide a double impact in strengthening both demand and supply sides of technical development.

Finally, patterns of demand emerge not just from developing skills and education in the population at large. They are affected also by government efforts at regulation. Government policy that strongly curtails private use of new technologies—excess taxation, insufficient access provision—can have adverse long-run consequences on growth in the economy.

References

- [1] Daron Acemoglu and Fabrizio Zilibotti. Productivity differences. Seminar Paper 660, IIES, University of Stockholm, Stockholm, November 1998.
- [2] Phillipe Aghion and Peter Howitt. A model of growth through creative destruction. *Econometrica*, 60(2):323–351, March 1992.
- [3] Kenneth J. Arrow. Economic welfare and the allocation of resources for inventions. In Richard R. Nelson, editor, *The Rate and Direction of Inventive Activity*. Princeton University Press and NBER, 1962.
- [4] W. Brian Arthur. *Increasing Returns and Path Dependence in the Economy*. The University of Michigan Press, Ann Arbor, 1994.
- [5] Partha Dasgupta. Patents, priority and imitation or, the economics of races and waiting games. *Economic Journal*, 98:66–80, March 1988.
- [6] Paul A. David. Clio and the economics of QWERTY. *American Economic Association Papers and Proceedings*, 75(2):332–337, May 1985.
- [7] Paul A. David. Knowledge, property, and the system dynamics of technological change. *Proceedings of the World Bank Annual Conference on Development Economics*, pages 215–248, March 1992.
- [8] Paul A. David. Intellectual property institutions and the panda’s thumb: Patents, copyrights, and trade secrets in economic theory and history. In M. B. Wallerstein, M. E. Moge, and R. A. Schoen, editors, *Global Dimensions of Intellectual Property Rights in Science and Technology*, chapter 2, pages 19–61. National Academy Press, 1993.
- [9] Thomas K. Derry and Trevor I. Williams. *A Short History of Technology: From the Earliest Times to A.D. 1900*. Oxford University Press, Oxford, 1960.
- [10] Stuart W. Elliott. Computer technology, human labor, and long-run economic growth. Working paper, Carnegie Mellon University, Pittsburgh, July 1998.

- [11] Gene M. Grossman and Elhanan Helpman. Quality ladders in the theory of growth. *Review of Economic Studies*, 58(1):43–61, January 1991.
- [12] Elhanan Helpman, editor. *General Purpose Technologies and Economic Growth*. MIT Press, Cambridge, Massachusetts, 1998.
- [13] Jon Holyoak and Paul Torremans. *Intellectual Property Law*. Butterworths, London, 1995.
- [14] Charles I. Jones. R&D-based models of economic growth. *Journal of Political Economy*, 103(3):759–784, August 1995.
- [15] Eric L. Jones. *Growth Recurring: Economic Change in World History*. Oxford University Press, Oxford, 1988.
- [16] Dale W. Jorgenson and Kevin Stiroh. Computers and growth. *Economics of Innovation and New Technology*, 3:295–316, 1995.
- [17] Kuk-Hwan Jung, Jeong-hun Oh, and Ilsoon Shin. The economic impact of information and communications technology in Korea. Working paper, National Computerization Agency, Korea, Seoul, November 1998.
- [18] Louise C. Keely and Danny Quah. Technology in growth. In Thijs Ruyter van Steveninck, editor, *Economic Growth and its Determinants*. Kluwer, 1999.
- [19] Adrienne Koch and William Peden, editors. *The Life and Selected Writings of Thomas Jefferson*. Random House, 1944.
- [20] Paul Krugman. *Geography and Trade*. MIT Press, 1991.
- [21] Per Krussel and José-Victor Rios-Rull. Vested interests in a positive theory of stagnation and growth. *Review of Economic Studies*, 63:301–329, April 1996.
- [22] David S. Landes. *The Wealth and Poverty of Nations*. Little, Brown and Company, London, 1998.
- [23] Glenn M. MacDonald. The economics of rising stars. *American Economic Review*, 78(1):155–166, March 1988.
- [24] William H. McNeill. *The Pursuit of Power*. University of Chicago, Chicago 60637, 1982.

- [25] Joel Mokyr. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press, Oxford, 1990.
- [26] Petri Niininen. Computers and economic growth in Finland. Working Paper 148, United Nations University/WIDER, Helsinki, August 1998.
- [27] Danny Quah. The invisible hand and the weightless economy. Occasional Paper 12, Centre for Economic Performance, LSE, London, May 1996.
- [28] Danny Quah. Twin peaks: Growth and convergence in models of distribution dynamics. *Economic Journal*, 106(437):1045–1055, July 1996.
- [29] Danny Quah. Empirics for growth and distribution: Polarization, stratification, and convergence clubs. *Journal of Economic Growth*, 2(1):27–59, March 1997.
- [30] Danny Quah. Increasingly weightless economies. *Bank of England Quarterly Bulletin*, 37(1):49–56, February 1997.
- [31] Danny Quah. Growth and wealth creation in the weightless, knowledge-based economy. In *Invitation pour la société de la connaissance. Aspects liés au management et à l'organisation du travail. Actus du colloque organisé le 7 Octobre 1998 par la Fondation de l'Entreprise*, Brussels, 1999. Fondation de l'Entreprise. (Published in 3 languages).
- [32] Danny Quah. Superstar knowledge-products in a model of growth. Working paper, LSE, London, January 1999.
- [33] Paul M. Romer. Endogenous technological change. *Journal of Political Economy*, 98(5, part 2):S71–S102, October 1990.
- [34] Sherwin Rosen. The economics of superstars. *American Economic Review*, 71(5):845–858, December 1981.
- [35] Avner Shaked and John Sutton. Relaxing price competition through product differentiation. *Review of Economic Studies*, 49:3–13, 1982.
- [36] Daniel E. Sichel. *The Computer Revolution: An Economic Perspective*. Brookings, Washington DC, 1997.

- [37] Poh-kam Wong. The contribution of information technology to the rapid economic growth of Singapore. Working paper, CMIT, National University of Singapore, Singapore, January 1999.
- [38] Brian D. Wright. The economics of invention incentives: Patents, prizes, and research contracts. *American Economic Review*, 73(4):691–707, September 1983.

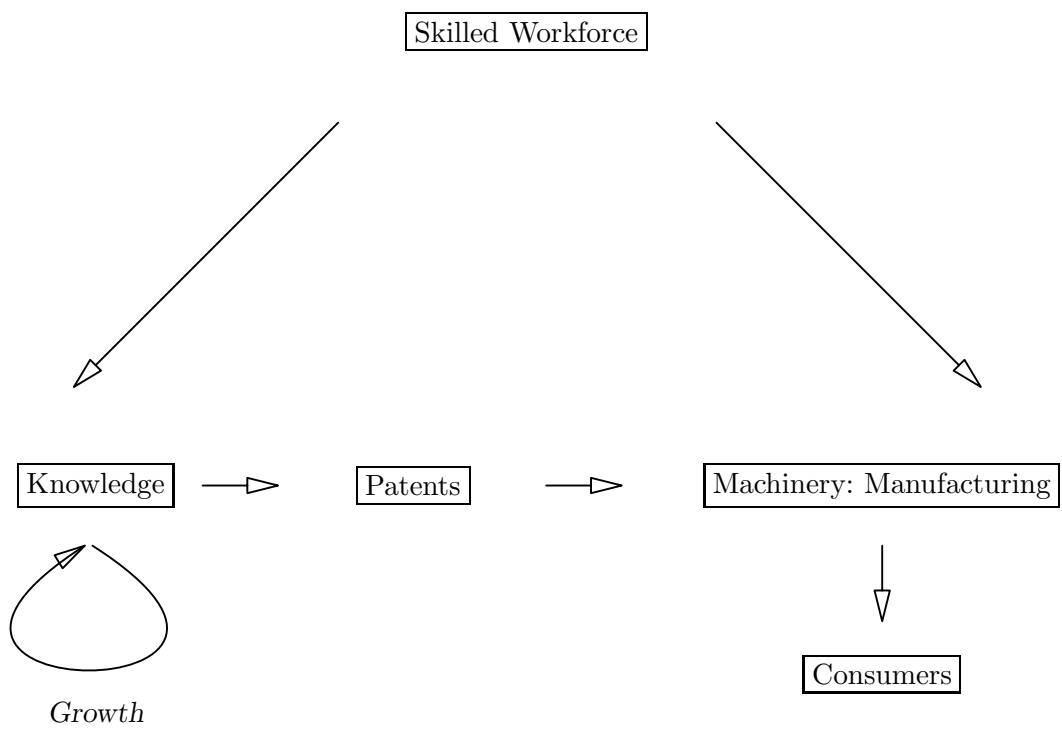


Fig. 1: Traditional models of knowledge and growth Patents and machinery intermediate between knowledge production and consumers




Fig. 2: Weightless economy Reduced distance between knowledge production and consumers

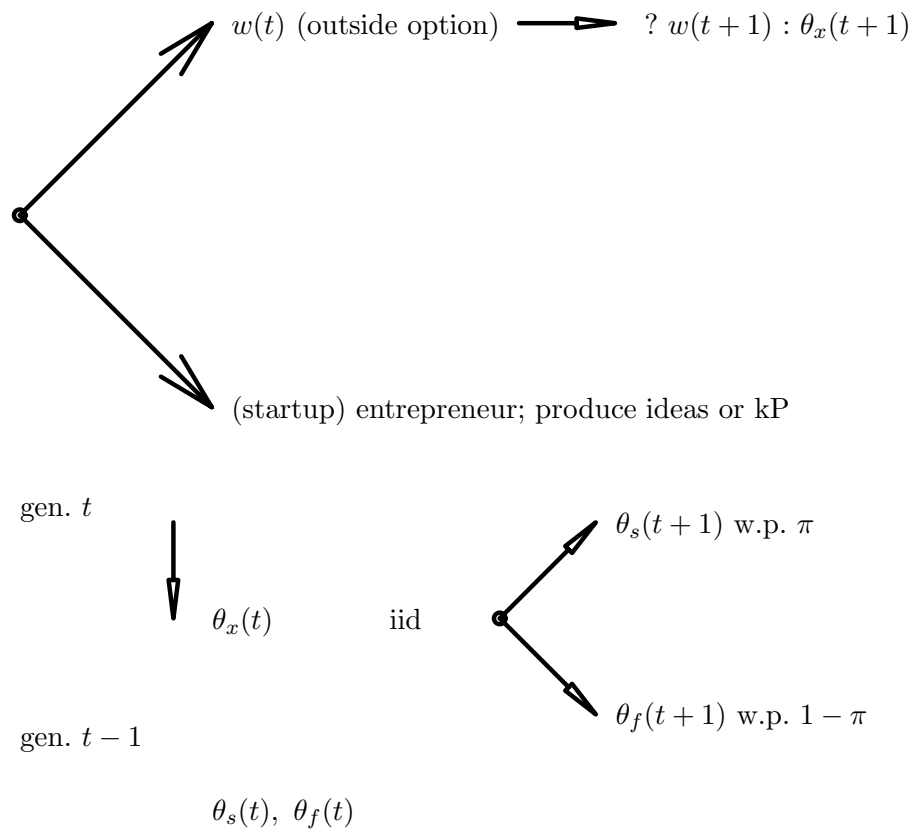


Fig. 3: Producer Demography and Decisions Environment and move sequence for generation t .




Fig. 4: Non-existent ν_{int} Only 2 kinds of consumption activity occur: One or the other kP (but not both) and no kP. The * and plain schedules can correspond to either s or x ; which each is depends on characteristics of the utility function and the possible range of variation in θ .




Fig. 5: Both intersections exist with $\nu_{max} < \nu_{int}$ Again, only 2 kinds of consumption activity occurring: One or the other kP (but not both) and no kP. The * and plain schedules can correspond to either s or x ; which each is depends on characteristics of the utility function and the possible range of variation in θ .




Fig. 6: Dispersed case with $\nu_{int} < \nu_{max}$ All 3 kinds of consumption activity occur. The * and plain schedules can correspond to either s or x ; which each is depends on characteristics of the utility function and the possible range of variation in θ .




Fig. 7: The learning society Intermediate ν 's consume experimentation.




Fig. 8: The conservative society Intermediate ν 's consume established successes.